

# Building Regular Registers with Rational Malicious Servers and Anonymous Clients – Extended Version

Antonella Del Pozzo<sup>1</sup>, Silvia Bonomi<sup>1</sup>, Riccardo Lazzeretti<sup>1</sup> and Roberto Baldoni<sup>1,2</sup>

<sup>1</sup> Research Center of Cyber Intelligence and Information Security (CIS), IT  
Dept. of Computer and System Sciences “Antonio Ruberti”, Sapienza Università di Roma, IT  
{delpozso, bonomi, lazzeretti, baldoni}@dis.uniroma1.it

<sup>2</sup> CINI Cybersecurity National Laboratory, Italy

**Abstract.** The paper addresses the problem of emulating a regular register in a synchronous distributed system where clients invoking `read()` and `write()` operations are anonymous while server processes maintaining the state of the register may be compromised by rational adversaries (i.e., a server might behave as *rational malicious Byzantine* process). We first model our problem as a Bayesian game between a client and a rational malicious server where the equilibrium depends on the decisions of the malicious server (behave correctly and not be detected by clients vs returning a wrong register value to clients with the risk of being detected and then excluded by the computation). We prove such equilibrium exists and finally we design a protocol implementing the regular register that forces the rational malicious server to behave correctly. **Keywords:** Regular Register, Rational Malicious Processes, Anonymity, Bayesian Game.

## 1 Introduction

To ensure high service availability, storage services are usually realized by replicating data at multiple locations and maintaining such data consistent. Thus, replicated servers represent today an attractive target for attackers that may try to compromise replicas correctness for different purposes, such as gaining access to protected data, interfering with the service provisioning (e.g. by delaying operations or by compromising the integrity of the service), reducing service availability with the final aim to damage the service provider (reducing its reputation or letting it pay for the violation of service level agreements), etc. A compromised replica is usually modeled through an arbitrary failure (i.e. a Byzantine failure) that is made transparent to clients by employing Byzantine Fault Tolerance (BFT) techniques. Common approaches to BFT are based on the deployment of a sufficiently large number of replicas to tolerate an estimated number  $f$  of compromised servers (i.e. BFT replication). However, this approach has a strong limitation: a smart adversary may be able to compromise more than  $f$  replicas in long executions and may get access to the entire system when the attack is sufficiently long. To overcome this issue, Sousa et al. designed the *proactive-reactive recovery* mechanism [21]. The basic idea is to periodically reconfigure the set of replicas to rejuvenate servers that may be under attack (proactive mode) and/or when a failure is detected (reactive mode). This approach is effective in long executions but requires a fine tuning of the replication parameters (upper bound  $f$  on the number of possible compromised replicas in a

given period, rejuvenation window, time required by the state transfer, etc...) and the presence of secure components in the system. In addition, it is extremely costly during good periods (i.e. periods of normal execution) as a high number of replicas must be deployed independently from their real need. In other words, the system pays the cost of an attack even if the attack never takes place.

In this paper, we want to investigate the possibility to implement a distributed shared variable (i.e. a register) without making any assumption on the knowledge of the number of possible compromised replicas, i.e. without relating the total number of replicas  $n$  to the number of possible compromised ones  $f$ . To overcome the impossibility result of [5,18], we assume that (i) clients preserve their privacy and do not disclose their identifiers while interacting with server replicas (i.e. anonymous clients) and (ii) at least one server is always alive and never compromised by the attacker. We first model our protocol as a game between two parties, a client and a rational malicious server (i.e. a server controlled by rational adversaries) where each rational malicious server gets benefit by two conflicting goals: (i) it wants to have continuous access to the current value of the register and, (ii) it wants to compromise the validity of the register returning a fake value to a client. However, if the rational malicious server tries to accomplish goal (ii) it could be detected by a client and it could be excluded from the computation, precluding it to achieve its first goal. We prove that, under some constraints, an equilibrium exists for such game. In addition, we design some distributed protocols implementing the register and reaching such equilibrium when rational malicious servers privilege goal (i) with respect to goal (ii). As a consequence, rational malicious servers return correct values to clients to avoid to be detected by clients and excluded by the computation and the register implementation is proved to be correct.

The rest of the paper is organized as follows: Section 2 discusses related works, Section 3 and Section 4 introduce respectively the system model and the problem statement. In Section 5 we model the problem as a Bayesian game and in Section 6 we provide a protocol matching the Bayesian Nash Equilibrium that works under some limited constraints, while in Section 7 we presents two variants of the protocol that relax the constraints, at the expense of some additional communications between the clients or protocol complexity increase. Finally, Section 8 presents a discussion and future work.

## 2 Related Work

Building a distributed storage able to resist arbitrary failures (i.e. Byzantine) is a widely investigated research topic. The Byzantine failure model captures the most general type of failure as no assumption is made on the behavior of faulty processes. Traditional solutions to build a Byzantine tolerant storage service can be divided into two categories: *replicated state machines* [19] and *Byzantine quorum systems* [5,16,17,18]. Both the approaches are based on the idea that the state of the storage is replicated among processes and the main difference is in the number of replicas involved simultaneously in the state maintenance protocol. Replicated state machines approach requires that every non-faulty replica receives every request and processes requests in the same order before returning to the client [19] (i.e. it assumes that processes are able to totally order requests and execute them according to such order). Given the upper bound on the num-

ber of failures  $f$ , the replicated state machine approach requires only  $2f + 1$  replicas in order to provide a correct register implementation. Otherwise, Byzantine quorum systems need just a sub-set of the replicas (i.e. *quorum*) to be involved simultaneously. The basic idea is that each operation is executed by a quorum and any two quorums must intersect (i.e. members of the quorum intersection act as witnesses for the correct execution of both the operations). Given the number of failures  $f$ , the quorum-based approach requires at least  $3f + 1$  replicas in order to provide a correct register implementation in a fully asynchronous system [18]. Let us note that, in both the approaches, the knowledge of the upper bound on faulty servers  $f$  is required to provide deterministic correctness guarantees. In this paper, we follow an orthogonal approach. We are going to consider a particular case of byzantine failures and we study the cost, in terms of number of honest servers, of building a distributed storage (i.e. a register) when clients are anonymous and have no information about the number of faulty servers (i.e. they do not know the bound  $f$ ). In particular, the byzantine processes here considered deviate from the protocol by following a strategy that brings them to optimize their own benefits (i.e., they are *rational*) and such strategy has the final aim to compromise the correctness of the storage (i.e., they are *malicious*). In [15], the authors presented Depot, a cloud storage system able to tolerate any number of Byzantine clients or servers, at the cost of a weak consistency semantics called *Fork-Join-Causal consistency* (i.e., a weak form of causal consistency).

Another different solution can rely on proactive secret sharing [25]. Secret Sharing [26] guarantees that a secret shared by a client among  $n$  parties (servers) cannot be obtained by an adversary corrupting no more than  $f$  servers ( $f$  imposed by the protocol). Moreover, if no more than  $f$  servers are Byzantines, the client can correctly recover the secret from the shares provided by any  $f + 1$  servers. Recent Proactive Secret Sharing protocols, e.g. [27], show that Secret sharing can be applied also to synchronous networks. Even if Proactive Secret Sharing can guarantee the privacy of the data (this is out of the scope of the paper) against up to  $n - 2$  passive adversaries, the solution has some limitations. First of all, clients are not able to verify whether the number of Byzantines exceeds  $t$  and hence understand if the message obtained is correct. Secondly, Secret Sharing protocols operating in a synchronous distributed system with Byzantines (active adversaries) correctly work with a small number of Byzantines and have high complexity ( $f < n/2 - 1$  and  $\mathcal{O}(n^4)$  in [27]).

In [3], the authors introduced the *BAR (Byzantine, Altruistic, Rational) model* to represent distributed systems with heterogeneous entities like peer-to-peer networks. This model allows to distinguish between Byzantine processes (arbitrarily deviating from the protocol, without any known strategy), altruistic processes (honestly following the protocol) and rational processes (may decide to follow or not the protocol, according to their individual utility). Under the BAR model, several problems have been investigated (e.g. reliable broadcast [7], data stream gossip [13], backup service through state machine replication [3]). Let us note that in the BAR model the utility of a process is measured through the cost sustained to run the protocol. In particular, each step of the algorithm (especially sending messages) has a cost and the objective of any rational process is to minimize its global cost. As a consequence, rational *selfish* processes deviate from the protocol just by skipping to send messages, if not properly encouraged

by some reward. In contrast with the BAR model, in this paper we consider malicious rational servers that can deviate from the protocol with different objectives, benefiting from preventing the correct protocol execution rather than from saving messages.

More recently, classical one-shot problems as leader election [1,2], renaming and consensus [2] have been studied under the assumption of rational agents (or rational processes). The authors provide algorithms implementing such basic building blocks, both for synchronous and asynchronous networks, under the so called *solution preference* assumption i.e., agents gain if the algorithm succeeds in its execution while they have zero profit if the algorithm fails. As a consequence, processes will not deviate from the algorithm if such deviation interferes with its correctness. Conversely, the model of rational malicious processes considered in this paper removes implicitly this assumption as they are governed by adversaries that get benefit when the algorithm fails while in [1,2] rational processes get benefit from the correct termination of the protocol (i.e. they are selfish according with the BAR model).

Finally, the model considered here can be seen as a particular case of BAR where rational servers take malicious actions, with the application similar to the one considered in [3]. However, in contrast to [3], we do not assume any trusted third party to identify users, we assume that clients are anonymous (e.g., they are connected through the Tor anonymous network [22]), and we investigate the impact of this assumption together with the rational model. To the best of our knowledge, this is the first paper that analyzes how the anonymity can help in managing rational malicious behaviors.

### 3 System Model

The distributed system is composed by a set of  $n$  servers implementing a distributed shared memory abstraction and by an arbitrary large but finite set of clients  $\mathcal{C}$ . Servers are fully identified (i.e. they have associated a unique identifier  $s_1, s_2 \dots s_n$ ) while clients are anonymous, i.e. they share the same identifier.

**Communication model and timing assumptions.** Processes can communicate only by exchanging messages through *reliable* communication primitives, i.e. messages are not created, duplicated or dropped. The system is synchronous in the following sense: all the communication primitives used to exchange messages guarantee a timely delivery property. In particular, we assume that clients communicate with servers through a *timely* reliable broadcast primitive (i.e., there exists an integer  $\delta$ , known by clients, such that if a client broadcasts a message  $m$  at time  $t$  and a server  $s_i$  delivers  $m$ , then all the servers  $s_j$  deliver  $m$  by time  $t + \delta$ ). Servers-client and client-client communications are done through “point-to-point” *anonymous timely* channels (a particular case of the communication model presented in [9] for the most general case of homonyms). Considering that clients are identified by the same identifier  $\ell$ , when a process sends a point-to-point message  $m$  to an identifier  $\ell$ , all the clients will deliver  $m$ . More formally, there exists an integer  $\delta' \leq \delta$ , known by processes, such that if  $s_i$  sends a message  $m$  to a client identified by an identifier  $\ell$  at time  $t$ , then all the clients identified by  $\ell$  receive  $m$  by time  $t + \delta'$ . We assume that channels are *authenticated* (“oral” model), i.e. when a process identified by  $j$  receives a message  $m$  from a process identified by  $i$ , then  $p_j$  knows that  $m$  has been generated by a process having identifier  $i$ .

**Failure model.** Servers are partitioned into two disjoint sub-sets: *honest* servers and *malicious* servers (*attackers*). Honest servers behave according to the protocol executed in the distributed system (discussed in Section 6) while malicious servers represent entities compromised by an adversary that may deviate from the protocol by dropping messages (omission failures), changing the content of a message, creating spurious messages, exchanging information outside the protocol, etc. Malicious servers are *rational*, i.e. they deviate from the protocol by following a strategy that aims at increasing their own benefit (usually performing actions that may prevent the correct execution of the protocol). We assume that rational malicious servers act independently, i.e. they do not form a coalition and each of them acts for its individual gain. Servers may also fail by crashing and we identify as *alive* the set of non crashed servers<sup>3</sup>. However, we assume that at least one honest alive server always exists in the distributed system.

## 4 Regular Registers

A register is a shared variable accessed by a set of processes, i.e. clients, through two operations, namely `read()` and `write()`. Informally, the `write()` operation updates the value stored in the shared variable while the `read()` obtains the value contained in the variable (i.e. the last written value). Every operation issued on a register is, generally, not instantaneous and it can be characterized by two events occurring at its boundary: an *invocation* event and a *reply* event. These events occur at two time instants (invocation time and reply time) according to the fictional global time.

An operation  $op$  is *complete* if both the invocation event and the reply event occur (i.e. the process executing the operation does not crash between the invocation and the reply). Contrary, an operation  $op$  is said to be *failed* if it is invoked by a process that crashes before the reply event occurs. According to these time instants, it is possible to state when two operations are concurrent with respect to the real time execution. For ease of presentation we assume the existence of a fictional global clock and the invocation time and response time of operations are defined with respect to this fictional clock. Given two operations  $op$  and  $op'$ , and their invocation event and reply event times ( $t_B(op)$  and  $t_B(op')$ ) and return times ( $t_E(op)$  and  $t_E(op')$ ), we say that  $op$  *precedes*  $op'$  ( $op \prec op'$ ) iff  $t_E(op) < t_B(op')$ . If  $op$  does not precede  $op'$  and  $op'$  does not precede  $op$ , then  $op$  and  $op'$  are *concurrent* ( $op || op'$ ). Given a `write( $v$ )` operation, the value  $v$  is said to be written when the operation is complete.

In case of concurrency while accessing the shared variable, the meaning of *last written value* becomes ambiguous. Depending on the semantics of the operations, three types of register have been defined by Lamport [14]: *safe*, *regular* and *atomic*. In this paper, we consider a regular register which is specified as follows:

- Termination: If an alive client invokes an operation, it eventually returns from that operation.
- Validity: A read operation returns the last value written before its invocation, or a value written by a write operation concurrent with it.

---

<sup>3</sup> Alive servers may be both honest or malicious.

Interestingly, safe, regular and atomic registers have the same computational power. This means that it is possible to implement a multi-writer/multi-reader atomic register from single-writer/single-reader safe registers. There are several papers in the literature discussing such transformations (e.g., [6,11,20,23,24] to cite a few). In this paper, we assume that the register is single writer in the sense that no two `write()` operations may be executed concurrently. However, any client in the system may issue a `write()` operation. This is not a limiting assumption as clients may use an access token to serialize their writes<sup>4</sup>. We will discuss in Section 8 how this assumption can be relaxed.

## 5 Modeling the Register protocol as a Game

In a distributed system where clients are completely disjoint from servers, it is possible to abstract any register protocol as a sequence of requests made by clients (e.g. a request to get the value or a request to update the value) and responses (or replies) provided by servers, plus some local computation. If all servers are honest, clients will always receive the expected replies and all replies will always provide the right information needed by the client to correctly terminate the protocol. Otherwise, a compromised server can, according to its strategy, omit to send a reply or can provide bad information to prevent the client from terminating correctly. In this case, in order to guarantee a correct execution, the client tries to detect such misbehavior, react and punish the server. Thus, a distributed protocol implementing a register in presence of rational malicious servers can be modeled as a two-party game between a client and each of the servers maintaining a copy of the register: the client wants to correctly access the register while the server wants to prevent the correct execution of a `read()` without being punished.

**Players.** The two players are respectively the client and the server. Each player can play with a different role: servers can be divided into *honest* servers and *malicious* servers while clients can be divided in those asking a *risky request* (i.e., clients able to detect misbehaviors and punish server<sup>5</sup>) and those asking for a *risk-less request* (i.e., clients unable to punish servers).


**Strategies.** Players' strategies are represented by all the possible actions that a process may take. Clients have just one strategy, identified by  $\mathcal{R}$ , that is *request information to servers*. Contrarily, servers have different strategies depending on their failure state:

- malicious servers have three possible strategies: (i)  $\mathcal{A}$ , i.e. *attack the client* by sending back wrong information (it can reply with a wrong value, with a wrong timestamp or both), (ii)  $\mathcal{NA}$ , i.e. *not attack the client* behaving according to the protocol and (iii)  $\mathcal{S}$ , i.e. *be silent* omitting the answer to client's requests;
- honest servers have just the  $\mathcal{NA}$  strategy.

Let us note that the game between a honest client and a honest server is trivial as they have just one strategy that is to follow the protocol. Thus, in the following we are going to skip this case and we will consider only the game between a client and a rational malicious server.

<sup>4</sup> Let us recall that we are in a synchronous system and the mutual exclusion problem can be easily solved also in presence of failures.

<sup>5</sup> Notice that the client ability to detect a server misbehaviors depends on the specific protocol.



**Fig. 1.** Extensive form of the game. Dashed line represents the unknown nature of requests from the risk point of view. Outcome pairs refer to client and server gains respectively.

**Utility functions and extensive form of the game.** Clients and servers have opposite utility functions. In particular:

- every client increases its utility when it is able to read a correct value from the register and it wants to maximize the number of successful `read()` operations;
- every server increases its utility when it succeeds to prevent the client from reading a correct value, while it loses when it is detected by the client and it is punished.

In the following, we will denote as  $G_c$  the gain obtained by the client when it succeeds in reading,  $G_s$  the gain obtained by the server when it succeeds in preventing the client from reading and as  $D_c$  the gain of the client when detecting the server and as  $D_s$  the loss of the server when it is detected. Such parameters are characteristic of every server and describe its behavior in terms of subjective gains/losses they are able to tolerate. Without loss of generality, we assume that  $G_c, G_s, D_c$  and  $D_s$  are all greater than 0, that all the servers have the same  $G_s$  and  $D_s$ <sup>6</sup> and that all the clients have the same  $G_c$  and  $D_c$ . Fig. 1 shows the extensive form of the game.

The game we are considering is a Bayesian game [10] as servers do not have knowledge about the client role but they can estimate the probability of receiving a risky request or a risk-less request i.e., they have a *belief* about the client role. We denote as  $\theta$  (with  $\theta \in [0, 1]$ ) the server belief of receiving a risky request (i.e. the client may detect that the server is misbehaving) and with  $1 - \theta$  the server belief of receiving a risk-less request (i.e. the client is not be able to detect that the server is misbehaving).

**Analysis of the Bayesian Game.** In the following, we are going to analyze the existence (if any) of a *Bayesian Nash Equilibrium* i.e., a Nash Equilibrium<sup>7</sup> computed by considering the players' belief.

Let us note that in our game, clients have just one strategy. Thus, the existence of the equilibrium depends only on the decisions taken by servers according to their utility parameters  $G_s, D_s$  and their belief about the nature of a request (i.e., its evaluation of  $\theta$ ). Let us now compute the expected gain  $E()$  of a server  $s_i$  while selecting strategies

<sup>6</sup> Let us note that if two servers have different values for  $G_s$  and  $D_s$ , the analysis shown in the following is simply repeated for each server.

<sup>7</sup> Let us recall that a Nash Equilibrium exists when each player selects a strategy and none of the players increases its utility by changing strategy.

$\mathcal{S}$ ,  $\mathcal{NA}$  and  $\mathcal{A}$ :

$$E(\mathcal{S}) = (-D_s \times (1 - \theta)) + (-D_s \times \theta) = -D_s \quad (1)$$

$$E(\mathcal{NA}) = ((1 - \theta) \times 0) + (\theta \times 0) = 0 \quad (2)$$

$$E(\mathcal{A}) = ((1 - \theta) \times G_s) - (\theta \times D_s) \quad (3)$$

**Lemma 1.** *The strategy  $\mathcal{S}$  is a dominated strategy.*

**Proof** A server  $s_i$  would choose to follow  $\mathcal{S}$  over  $\mathcal{NA}$  or  $\mathcal{A}$  if (i)  $E(\mathcal{S}) > E(\mathcal{A})$  or  $E(\mathcal{S}) > E(\mathcal{NA})$ . However, considering that both  $D_s$  and  $G_s$  are greater than 0, from equations (1) – (3) we will have that  $s_i$  will never choose to play  $\mathcal{S}$ .  $\square_{\text{Lemma 1}}$

It follows that servers have no gain in playing  $\mathcal{S}$ , whatever the other player does (cf. Lemma 1). In fact, there would be no increment of their utility by playing  $\mathcal{S}$  and then we will not consider such strategy anymore.

Let us note that a server  $s_i$  would prefer to play  $\mathcal{NA}$  (i.e., to behave honestly) with respect to  $\mathcal{A}$  (i.e., to deviate from the protocol) when  $E(\mathcal{NA}) > E(\mathcal{A})$ . Combining equations (3) and (2) we have that a  $s_i$  would prefer to play  $\mathcal{NA}$  when

$$\frac{G_s}{(G_s + D_s)} > \theta. \quad (4)$$

The parameters  $G_s$  and  $D_s$  are strictly dependent on the attacker's profile (i.e., an attacker for which it is more important to stay in the system rather than subvert it or vice versa), thus we can not directly work on them. In the remaining part of the work we propose protocols to tune the  $\theta$  parameter in such a way that the inequality (4) holds. To this purpose, we derive the following Lemmas:

**Lemma 2.** *Let  $s_i$  be a rational malicious server. If  $D_s < G_s$  and  $\theta < \frac{1}{2}$  then the best response of  $s_i$  is to play strategy  $\mathcal{A}$  (i.e.  $\mathcal{NA}$  is a dominated strategy).*

**Proof** Equation (4) can be rewritten as

$$\frac{1}{1 + \alpha} > \theta$$

where  $\alpha = \frac{D_s}{G_s}$ . Considering that  $G_s > D_s$  it follows that  $\alpha \in (0, 1)$ . Note that  $s_i$  would prefer to play  $\mathcal{A}$  each time that inequality (4) is satisfied and that  $\theta$  is upper bounded by  $\frac{1}{2}$ , it follows that  $s_i$  will prefer to play  $\mathcal{A}$  for any  $\theta \in [0, \frac{1}{2})$  and the claim follows.  $\square_{\text{Lemma 2}}$

**Lemma 3.** *Let  $s_i$  be a rational malicious server. If  $D_s > G_s$  and  $\theta \geq \frac{1}{2}$  then the best response of  $s_i$  is to never play strategy  $\mathcal{A}$  (i.e.  $\mathcal{NA}$  is a dominant strategy).*

**Proof** Equation (4) can be rewritten as

$$\frac{1}{1 + \alpha} > \theta$$



Note that  $G_s < D_s$  and then  $\alpha \in (1, \infty)$ . Considering that  $s_i$  would prefer to play  $\mathcal{A}$  each time that inequality (4) is satisfied and that  $\theta$  is lower bounded by  $\frac{1}{2}$ , it follows that  $s_i$  will never prefer to play  $\mathcal{A}$  for any  $\theta \in [\frac{1}{2}, 1]$  and the claim follows.

□ *Lemma 3*

## 6 A Protocol $\mathcal{P}$ for a Regular Register when $D_s \gg G_s$

In this section, we propose a protocol  $\mathcal{P}$  implementing a regular register in a synchronous distributed system with anonymous clients and up to  $n - 1$  malicious rational servers. The protocol works under the assumption that the server loss  $D_s$  in case of detection is much higher than its gain  $G_s$  obtained when the client fails during a read (i.e.  $D_s \gg G_s$ <sup>8</sup>). This assumption models a situation where the attacker is much more interested in having access to data stored in the register and occasionally interfere with the server rather than causing a reduction of the availability (e.g., no termination or validity violation). We will relax this assumption to the simple case  $D_s > G_s$  in the next section extending  $\mathcal{P}$  in two different ways.

Our protocol  $\mathcal{P}$  follows the classical quorum-based approach. When a client wants to write, it sends the new value together with its timestamp to servers and waits for acknowledgments. Similarly, when it wants to read, it asks for values and corresponding timestamps and then it tries to select a value among the received ones. Let us note that, due to the absence of knowledge on the upper bound of malicious processes, it could be impossible for a reader to select a value among those reported by servers and, in addition, the reader may be unable to distinguish well behaving servers from malicious ones. To overcome this issue we leverage on the following observation: the last client  $c_w$  writing a value  $v$  is able to recognize such value while reading after its write (as long as no other updates have been performed). This makes the writer  $c_w$  the only one able to understand which server  $s_i$  is reporting a wrong value  $v_i \neq v$ , detect it as malicious and punish it by excluding  $s_i$  from the computation. Thus, the basic idea behind the protocol is to exploit the synchrony of the system and the anonymity of clients to make the writer indistinguishable from readers and “force” malicious servers to behave correctly.

Let us note that anonymity itself is not enough to make the writer indistinguishable from other clients. In fact, if we consider a naive solution where we add anonymity to a register implementation (e.g., to the one given by Attiya, Bar-Noy and Dolev [4]), we have that servers may exploit the synchrony of the channels to estimate when the end of the write operation occurs and to infer whether a read request may arrive from the writer or from a different client (e.g., when it is received too close to a write request and before the expected end of the write). To this aim, we added in the `write()` operation implementation some *dummy* read requests. These messages are actually needed to generate message patterns that make impossible to servers to distinguish messages coming from the writer from messages arriving from a different client. As a consequence, received a read request, a server  $s_i$  is not able to distinguish if such request is risky (i.e. it comes from the writer) or is risk-less (i.e. it comes from a generic client).

<sup>8</sup> More precisely,  $\mathcal{P}$  works when  $D_s > cG_s$  where  $c$  is the estimated number of clients in the system.

In addition, we added a detection procedure that is executed both during `read()` and `write()` operations by any client. In particular, such procedure checks that every server answered to a request and that the reported information are “coherent” with its knowledge (e.g., timestamps are not too old or too new). The detection is done first locally, by exploiting the information that clients collect during the protocol execution, and then, when a client detects a server  $s_j$ , it disseminates its detection so that the malicious server is permanently removed from the computation (collaborative detection).

Finally, the timestamp used to label a new written value is updated by leveraging acknowledgments sent by servers at the end of the preceding `write()` operation. In particular, during each `write()` operation, servers must acknowledge the write of the value by sending back the corresponding timestamp. This is done on the anonymous channels that deliver such message to all the clients that will update their local timestamp accordingly. As a consequence, any rational server is inhibited from deviating from the protocol, unless it accepts the high risk to be detected as faulty and removed from the system.

In the following, we provide a detailed description of the protocol  $\mathcal{P}$  shown in Figures 2-4.

**The `read()` operation (Fig. 2).** When a client wants to read, it first checks if the *last\_ts* variable is still equal to 0. If so, then there is no `write()` operation terminated before the invocation of the `read()` and the client returns the default value  $\perp$  (line 04, Fig. 2(a)). Otherwise,  $c_i$  queries the servers to get the last value of the register by sending a `READ()` message (line 06, Fig. 2(a)) and remains waiting for  $2\delta$  times, i.e. the maximum round trip message delay (line 07, Fig. 2(a)).

When a server  $s_i$  delivers a `READ()` message, the *reading<sub>i</sub>* counter is increased by one and then  $s_i$  sends a `REPLY(< i, tsi, vali, old_tsi, old_vali >)` message containing the current and old values and timestamp stored locally (lines 03 - 04, Fig. 2(b)).

When the reading client delivers a `REPLY(< j, ts, val, ots, ov >)` message, it stores locally the reply in two tuples containing respectively the current and the old triples with server id, timestamp and corresponding value (lines 24 - 25, Fig. 2(a)). When the reader client is unblocked from the wait statement, it checks if there exists a pair  $\langle ts, val \rangle$  in the *replies* set that has been reported by all servers it believes honest (line 08, Fig. 2(a)) and, in this case, it sends a `READ_ACK()` message (line 09, Fig. 2(a)) and it returns the corresponding value (line 10, Fig. 2(a)). Received the `READ_ACK()` message, a server  $s_i$  just decreases by one its *reading<sub>i</sub>* counter (line 05, Fig. 2(b)). Otherwise, a `write()` operation may be in progress. To check if it is the case, the client keeps waiting for other  $\delta$  time units and then checks again if a good value exists (lines 11 - 12, Fig. 2(a)). If, after this period, the value is not yet found, it means that some of the servers behaved maliciously. Therefore, the client executes the `detection()` procedure to understand who is misbehaving (cfr. Fig. 4). Let us note that such procedure cleans up the set of honest servers when they are detected to be malicious. Therefore, after the execution of the procedure, the reader checks for the last time if a good value exists in its *replies* set and, if so, it returns such value (line 18, Fig. 2(a)); otherwise the special value abort is returned (line 19, Fig. 2(a)). In any case, a `READ_ACK()` is sent to block the forwarding of new values at the server side (line 16, Fig. 2(a)).

```

Init:
(01)  $replies \leftarrow \emptyset; my\_last\_val \leftarrow \perp; my\_last\_ts \leftarrow 0; last\_ts \leftarrow 0;$ 
(02)  $ack \leftarrow \emptyset; honest \leftarrow \{s_1, s_2 \dots s_n\}; writing \leftarrow false;$ 



---


operation read():
(03) if ( $last\_ts = 0$ )
(04)   then return  $\perp$ ;
(05)   else  $replies \leftarrow \emptyset;$ 
(06)     broadcast READ();
(07)     wait ( $2\delta$ );
(08)     if ( $\forall s_i \in honest, \exists \langle -, ts, val \rangle \in replies$ )
(09)       then broadcast READACK();
(10)       return  $val$ ;
(11)     else wait ( $\delta$ );
(12)       if ( $\forall s_i \in honest, \exists \langle -, ts, val \rangle \in replies$ )
(13)         then broadcast READACK();
(14)         return  $val$ ;
(15)       else execute detection( $replies_i, R$ )
(16)         broadcast READACK();
(17)         if ( $\forall s_i \in honest, \exists \langle -, ts, val \rangle \in replies$ )
(18)           then return  $val$ ;
(19)           else abort ;
(20)         endif
(21)       endif
(22)     endif
(23)   endif



---


when  $REPLY(\langle j, ts, v, ots, ov \rangle)$  is delivered:
(24)  $replies \leftarrow replies \cup \{\langle j, ts, v \rangle\};$ 
(25)  $replies \leftarrow replies \cup \{\langle j, ots, ov \rangle\};$ 



---


when DETECTED( $s_j$ ) is delivered:
(26)  $honest \leftarrow honest \setminus \{s_j\};$ 

```

(a) Client Protocol

```

Init:
(01)  $val_i \leftarrow \emptyset; ts_i \leftarrow 0;$ 
(02)  $old\_val_i \leftarrow \perp; old\_ts_i \leftarrow 0; reading_i \leftarrow 0;$ 



---


when READ() is delivered:
(03)  $reading_i \leftarrow reading_i + 1;$ 
(04) send  $REPLY(\langle i, ts_i, val_i, old\_ts_i, old\_val_i \rangle);$ 



---


when READACK() is delivered:
(05)  $reading_i \leftarrow reading_i - 1;$ 

```

(b) Server Protocol

**Fig. 2.** The read() protocol for a synchronous system.

**The write() operation (Fig. 3).** When a client wants to write, it first sets its *writing* flag to true, stores locally the value and the corresponding timestamp, obtained incrementing by one the current timestamp stored in *last\_ts* variable (lines 01 - 02, Fig. 3(a)), sends a WRITE() message to servers, containing the value to be written and the corresponding timestamp (line 03, Fig. 3(a)), and remains waiting for  $\delta$  time units. When a server  $s_i$  delivers a WRITE( $v, ts$ ) message, it checks if the received timestamp is greater than the one stored in the  $ts_i$  variable. If so,  $s_i$  updates its local variables keeping the current value and timestamp as old and storing the received ones as current (lines 02 - 05, Fig. 3(b)). Contrarily,  $s_i$  checks if the timestamp is the same stored locally in  $ts_i$ . If this happens, it just adds the new value to the set  $val_i$  (line 06, Fig.



**Fig. 3.** write() protocol for a synchronous system.

3(b)). In any case,  $s_i$  sends back an ACK() message with the received timestamp (lines 08, Fig. 3(b)) and forwards the new value if some read() operation is in progress (lines 09, Fig. 3(b)). Delivering an ACK() message, the writer client checks if the timestamp is greater equal than its  $my\_last\_ts$  and, if so, it adds a tuple  $\langle j, ts, - \rangle$  to its  $ack$  set (line 16, Fig. 3(a)).

When the writer is unblocked from the wait statement, it sends a READ() message, waits for  $\delta$  time units and sends another READ() message (lines 06 - 08, Fig. 3(a)). This message has two main objectives: (i) create a message pattern that makes impossible to malicious servers to distinguish a real reader from the writer and (ii) collect values to detect misbehaving servers. In this way, a rational malicious server, that aims at remaining in the system, is inhibited from misbehaving as it could be detected from the writer and removed from the computation. The writer, in fact, executes the detection() procedure both on the  $ack$  set and on the  $replies$  set collected during the write() (lines 09 - 11, Fig. 3(a)). Finally, the writer sends two READ\_ACK() messages to block the forwarding of replies, resets its  $writing$  flag to false and returns from the operation (lines 12 - 15, Fig. 3(a)).

```

procedure detection(replies_set, set_type):
(01)  $S = \{j | \exists <j, -, -> \in \text{replies\_set}\}$ ;
(02) if (honest  $\not\subseteq S$ )
(03)   then for each  $s_j \in (\text{honest}_i \setminus S)$  do
(04)     trigger detect( $s_j$ );
(05)      $\text{honest}_i \leftarrow \text{honest}_i \setminus \{s_j\}$ ;
(06)     broadcast DETECTED( $s_j$ );
(07)   endFor
(08) endif
(09) if (set_type = R)
(10)   then if (writing)
(11)     then  $R = \{j | \exists <j, \text{my\_last\_val}, \text{my\_last\_ts}> \in \text{replies\_set}\}$ ;
(12)     if (honest  $\not\subseteq R$ )
(13)       then for each  $s_j \in (\text{honest}_i \setminus R)$  do
(14)         trigger detect( $s_j$ );
(15)          $\text{honest}_i \leftarrow \text{honest}_i \setminus \{s_j\}$ ;
(16)         broadcast DETECTED( $s_j$ );
(17)       endFor
(18)     endif
(19)   else for each  $<j, ts, -> \in \text{replies\_set}$  such that  $ts < \text{last\_ts} - 1$  do
(20)     trigger detect( $s_j$ );
(21)      $\text{honest} \leftarrow \text{honest} \setminus \{s_j\}$ ;
(22)     broadcast DETECTED( $s_j$ );
(23)   endFor
(24)   for each  $<j, ts, \text{val}> \in \text{replies\_set}$  such that  $ts = \text{my\_last\_ts}$  do
(25)      $D_i = \{v | (\exists <j, ts, \text{val}> \in \text{replies\_set}) \wedge (ts = \text{my\_last\_ts})\}$ ;
(26)     if ( $(\text{my\_last\_val} \neq \perp) \wedge (\text{my\_last\_ts} = \text{last\_ts}) \wedge (\text{last\_val} \notin D_i)$ )
(27)       then trigger detect( $s_j$ );
(28)        $\text{honest} \leftarrow \text{honest} \setminus \{s_j\}$ ;
(29)       broadcast DETECTED( $s_j$ );
(30)     endif
(31)   endFor
(32)   for each  $<j, ts, \text{val}> \in \text{replies\_set}$  such that  $ts > \text{last\_ts} + 1$  do
(33)     trigger detect( $s_j$ );
(34)      $\text{honest}_i \leftarrow \text{honest}_i \setminus \{s_j\}$ ;
(35)     broadcast DETECTED( $s_j$ );
(36)   endFor
(37) endif
(38) else for each  $<j, ts, -> \in \text{replies\_set}$  such that  $ts \neq \text{my\_last\_ts}$  do
(39)   trigger detect( $s_j$ );
(40)    $\text{honest} \leftarrow \text{honest} \setminus \{s_j\}$ ;
(41)   broadcast DETECTED( $s_j$ );
(42) endFor
(43) endif.

```

**Fig. 4.** detection() function invoked by an anonymous client for a synchronous system.

Let us note that, the execution of a write() operation triggers the update of the *last\_ts* variable at any client. This happens when in the *ack* set there exists a timestamp reported by any honest server (lines 17 - 18, Fig. 3(a)).

**The detection() procedure (Fig 4).** This procedure is used by clients to detect servers misbehaviors during the execution of read() and write() operations. It takes as parameter a set (that can be the *replies* set or the *ack* set) and a flag that identifies the type of the set (i.e. *A* for *ack*, *R* for *replies*). In both cases, the client checks if it has received at least one message from any server it saw honest and detects as faulty all the servers omitting a message (lines 01 - 08).

If the set to be checked is a set of ACK() messages, the client (writer) just checks if some server  $s_j$  acknowledged a timestamp that is different from the one it is using in

the current `write()` and, if so,  $s_j$  is detected as malicious (lines 38 - 42). Otherwise, if the set is the *replies* set (flagged as  $R$ ), the client checks if it is running the procedure while it is writing or reading (line 10). If the client is writing, it just updated the state of the register. Thus, the writer checks that all servers sent back the pair  $\langle v, ts \rangle$  corresponding to the one stored locally in the variables *myLastVal* and *myLastTs*. If someone reported a bad value or timestamp, it is detected as misbehaving (lines 11 - 18). If the client is reading, it is able to detect servers sending back timestamps that are too old (lines 19 - 23) or too new to be correct (lines 32 - 36) or servers sending back the right timestamp but with a wrong value (lines 24 - 31).

## 6.1 Correctness Proofs for $\mathcal{P}$

In the following we prove that the protocol presented in Fig. 2 - 4 terminates (Lemma 4, Lemma 5 and Theorem 1), and demonstrate some properties of the timestamp mechanism used to label `write()` operations. In particular, we prove that the protocol ensures the increasing monotonic order of timestamps (Lemma 8) and the consistency of the variable storing the last used timestamp (Lemma 7). In Lemma 9 we prove that the last written value persists locally at each honest server, while Theorem 2 proves that if servers behave honestly then the protocol emulates a regular register. Then we prove the accuracy of the detection function (Lemma 11, Lemma 12 and Theorem 3) and that the proposed protocol  $\mathcal{P}$  emulates a regular register under specific conditions (Theorem 4).

**Lemma 4.** *Let  $c_\ell$  be an anonymous client invoking a `write()` operation. If  $c_\ell$  executes the protocol in Fig. 3 then it eventually returns from the `write()` operation.*

**Proof** The proof simply follows by observing that in the `write()` operation code (e.g. Fig. 3(a)) the return event happens after three `wait()` statement. Thus, considering that  $c_\ell$  is honest, it will be unblocked, from the last `wait()` statement,  $3\delta$  time units after the `write()` invocation and the claim follows.  $\square_{\text{Lemma 4}}$

**Lemma 5.** *Let  $c_\ell$  be an anonymous client invoking a `read()` operation. If  $c_\ell$  executes the protocol in Fig. 2 then it eventually returns from the `read()` operation.*

**Proof** The proof simply follows by observing that in the `read()` operation code a return event is defined in every branch of the code and it only depends on `wait()` statements. Thus, considering that  $c_\ell$  is honest, it will be unblocked in a finite time after the `read()` invocation and the claim follows.  $\square_{\text{Lemma 5}}$

**Theorem 1 (Termination).** *Let  $c_\ell$  be an anonymous client invoking an operation  $op$ . If  $c_\ell$  executes the protocol in Fig. 2 - 4 then it eventually returns from  $op$ .*

**Proof** The proof directly follows from Lemma 4 and Lemma 5.  $\square_{\text{Theorem 1}}$

**Lemma 6.** *Let  $op$  be a  $\text{write}()$  operation and let  $ts$  be the timestamp associated by the writer to  $op$ . If a malicious server  $s_i$  deviates from the protocol by omitting the  $\text{WRITE\_ACK}(ts)$  message or by sending a  $\text{WRITE\_ACK}(ts')$  message (with  $ts' \neq ts$ ), then  $s_i$  will be detected as malicious by any client.*

**Proof** The  $\text{detection}()$  function is executed by the writer client on the  $ack$  set at line 09, Fig. 3(a). The  $ack$  set is emptied at the beginning of every  $\text{write}()$  operation and it is filled-in by the writer when it deliver  $\text{WRITE\_ACK}(ts')$  messages. Such messages are sent by servers when delivering a  $\text{WRITE}(< ts, val >)$  message sent by the writer at the beginning of the operation. The proof simply follows by considering that the writer client knows the real value of the timestamp associated to the write and stored in its  $my\_last\_ts$  local variable (line 02, Fig. 3(a)). Thus, when the writer executes the  $\text{detection}()$  function the writer checks (i) if it has received a  $\text{WRITE\_ACK}()$  message from any servers it sees (line 01-07, Fig. 4) and (ii) if all the alive servers acknowledge the right timestamp (line 38-42, Fig. 4). Considering that the communications are timely and the detection happens  $2\delta$  time units after the broadcast of the  $\text{WRITE}(< ts, val >)$  message (i.e. after the maximum round trip delay), if some servers do not answer they are detected as malicious as they omitted to answer. In the second case, since the writer know its timestamp and channel are authenticated, it is able to detect as malicious the server answering with a different timestamp. Finally, considering that (i) the writer notifies to all the other clients its detections, (ii) such detections are done  $\delta$  time units before the end of the write and (iii) clients notifications delay is also bounded by  $\delta$ , it follows that at the end of the write any client detected the malicious servers and the claim follows.  $\square_{\text{Lemma 6}}$

**Lemma 7.** *At the end of every  $\text{write}()$  operation any client stores in its  $last\_ts$  variable the same timestamp.*

**Proof** Every client initializes its  $last\_ts$  variable to 0 during the init phase (line 01, Fig. 2(a)). Such variable is updated at line 17, Fig. 3(a) when the client stores in its  $ack$  set a timestamp  $ts'$  that has been acknowledged by any alive server and that is greater than the previous one (to preserve the monotonically increasing order of timestamps). Then, we just need to prove that if a client updates its  $last\_ts$  variable then all the clients will update it as well.

Considering that (i)  $\text{WRITE\_ACK}(ts')$  messages are sent by servers when a  $\text{WRITE}(< val, ts' >)$  message is delivered, (ii)  $\text{WRITE}(< val, ts' >)$  messages are sent to all the servers and (iii)  $\text{WRITE\_ACK}(ts')$  messages are sent through point-to-point anonymous channels, we have that all clients receive  $\text{WRITE\_ACK}(ts')$  messages from the same set of servers. In addition, due to Lemma 6, we have that the set of alive processes is shared by every client and the claim follows.  $\square_{\text{Lemma 7}}$

**Lemma 8.** *Let  $op$  and  $op'$  be two  $\text{write}()$  operations such that  $op \prec op'$ . Let  $ts$  and  $ts'$  be respectively the timestamp associated to  $op$  and to  $op'$ , then  $ts < ts'$ .*

**Proof** The proof simply follows by Lemma 7 and considering that the timestamp associated to a  $\text{write}()$  operation is computed by incrementing the  $\text{last\_ts}$  variable by one.

□<sub>Lemma 8</sub>

**Corollary 1.** *If there not exists two concurrent  $\text{write}()$  operations, for any pair of  $\text{write}()$   $w_1, w_2$  such that  $w_1 \prec w_2$  and there not exists any  $w_3$  such that  $w_1 \prec w_3 \prec w_2$  then the timestamp  $ts_2$  associated to  $w_1$  is equal to  $ts_1 + 1$ .*

**Proof** The proof simply follows by considering that timestamps are computed by incrementing the  $\text{last\_ts}$  variable and it is updated at most once during each  $\text{write}()$  operation.

□<sub>Corollary 1</sub>

**Lemma 9.** *At the end of a  $\text{write}(v)$  operation, every server  $s_i$  behaving honestly stores the value  $v$  in its  $\text{val}_i$  local variable.*

**Proof** Every server  $s_i$  updates its  $\text{value}_i$  variable in line 05 or line 06, Fig. 3(b). In particular, this happens when the timestamp attached to the  $\text{WRITE}()$  message and associated to the  $\text{write}()$  operation is greater equal than the one stored in  $ts_i$ . Due to Lemma 8, timestamps follows a monotonically increasing order and thus every  $\text{write}()$  operation will have a timestamp that is greater or equal than the one previously stored. As a consequence, when delivering a  $\text{WRITE}(< \text{val}, ts' >)$  message, any server  $s_i$  will always execute line 05 or 06, Fig. 3(b) storing locally the new value and the claim follows.

□<sub>Lemma 9</sub>

To the ease of presentation, let us assume that the default value  $\perp$  is written by a fictional instantaneous  $\text{write}(\perp)$  operation preceding every operation  $op$  executed by clients and let us define a valid value as follows:

**Definition 1.** *Let  $v$  be the value returned by a  $\text{read}()$  operation  $op$  invoked on the regular register.  $v$  is said to be valid if*

- (i) *it is the value written by the last  $\text{write}()$  operation terminated before  $op$  or*
- (ii) *it is the value written by a  $\text{write}()$  operation concurrent with  $op$ .*

**Lemma 10.** *If all the servers behave honestly (i.e. they follow the protocol presented in Fig. 2(b) - 3(b)) then any  $\text{read}()$  operation returns a valid value.*

**Proof** Let us suppose by contradiction that all servers follow the protocol and that there exists a  $\text{read}()$  operation  $op$  that returns a value  $v$  that is not valid.

Let  $v_{old}$  be the value written by the last  $\text{write}()$  terminated before the invocation of  $op$ .

If  $v$  is not valid, it means that  $v$  is different from  $v_{old}$  and from the value  $v'$  written by a concurrent  $\text{write}(v')$ , if it exists.

**Case 1 - No  $\text{write}()$  operation is concurrent with  $op$ .** Due to Lemma 9, at time  $t$  when  $op$  is invoked, any server will store locally in their  $\text{value}_i$  variable the value  $v_{old}$ , together with its timestamp. Since the value stored in the register is updated only when a  $\text{WRITE}()$  message is delivered (line 05 or 06, Fig. 3(b)), and this happens only when



a  $\text{write}()$  operation is triggered (line 03, Fig. 3(a)), we have that, if no  $\text{write}()$  operation is concurrent with  $op$ ,  $v\_old$ , together with its timestamp, will be stored and not updated by any server during the whole execution of  $op$ . Considering that any server  $s_i$  behaves honestly, while delivering a  $\text{READ}()$  message,  $s_i$  will answer by sending a  $\text{REPLY}(< i, ts_i, value_i >)$  containing the same value and the same timestamp to the reader (line 04, Fig. 2(b)). Thus, at time  $t + 2\delta$  the  $replies$  set will contains  $n$  tuple  $< v\_old, ts >$  and the condition at line 08, Fig. 2(a) holds terminating the  $\text{read}()$  operation with  $v\_old$  and we have a contradiction as it is a valid value.

**Case 2 - There exists a  $\text{write}(v')$  operation  $op'$  concurrent with  $op$ .** Without loss of generality, let  $x$  be the timestamp associated to the value  $v\_old$  written by the last terminated  $\text{write}()$  operation preceding  $op$ . Let us denote by  $op'$  the  $\text{write}(v')$  operation concurrent with the  $\text{read}()$  operation  $op$ . Let us note that, according to the protocol in Fig. 2(a), while executing the  $\text{read}()$  operation  $op$ , the reader client will inquiry servers to get the value of the register together with its timestamp. Considering that, by assumption, any alive server  $s_i$  behaves honestly, it follows that delivering a  $\text{READ}()$  message  $s_i$  will answer by sending back a  $\text{REPLY}()$  message containing both the current and the old value and timestamp. Such values are modified only at line 05 or 06, Fig. 3(b) when a server  $s_i$  deliver a  $\text{WRITE}()$  message.

Let  $t_B(op')$  and  $t_B(op)$  the time at which respectively  $op'$  and  $op$  are invoked and let  $t_E(op')$  and  $t_E(op)$  be respectively the return time of  $op'$  and  $op$ .

Let us consider the following cases:

– **Case 2.1 -  $t_B(op') + \delta < t_B(op) < t_B(op') + 2\delta$ .**

At the beginning of the  $\text{write}()$  (i.e. at time  $t_B(op')$ ), the writer client sends a  $\text{WRITE}()$  message (line 03, Fig. 3(a)) that will be delivered by any alive server by time  $t_B(op') + \delta$ . When a server  $s_i$  receives a  $\text{WRITE}()$  message, it will update its  $val_i$  variable to  $v'$  and its  $ts_i$  to the current timestamp<sup>9</sup>. It follow that from time  $t_B(op') + \delta$  any alive server will store locally the value  $v'$  written by  $op'$ . Let us note that, since  $op'$  lasts until time  $t_E(op') = t_B(op') + 3\delta$  and, by assumption, there not exist concurrent  $\text{write}()$  operations, such variables will not be updated anymore before time  $t_B(op') + 3\delta$ .

When an alive server  $s_i$  delivers a  $\text{READ}()$  message (between time  $t_B(op') + \delta$  and  $t_B(op') + 2\delta$ ), it executes line 04, Fig. 2(b) sending to the reader both the current and the old values and timestamp through a  $\text{REPLY}()$  message. Delivering such  $\text{REPLY}()$  message, the reader will store the values in its  $replies$  local variable and waits until time  $t_B(op) + 2\delta$  before checking the content of such variable. Considering that, by assumption, all the servers behave honestly, they will send the content of their  $val_i$  and  $ts_i$  variable without changing them. As a consequence, the reader will store locally in its  $replies$  set the same pair  $< v', ts >$  from any alive server. Thus, at time  $t_B(op) + 2\delta \leq t_B(op') + 3\delta$ , evaluating the condition at line 08, Fig. 2(a) the reader will select and return  $v'$  and we have a contradiction.

– **Case 2.2 -  $t_B(op) < t_B(op') + \delta$ .**

In this case, the  $\text{WRITE}()$  message is concurrent (wrt. the happened before relation)

<sup>9</sup> Let us note that such variable will be always updated, as in line 05 or 06, Fig. 3(b), due to Lemma 8.

with the  $\text{READ}()$  message. As a consequence, we may have that a server  $s_i$  delivering the  $\text{READ}()$  message before than the  $\text{WRITE}()$  message and a server  $s_j$  delivering the  $\text{WRITE}()$  message before than the  $\text{READ}()$  message. As a consequence,  $s_i$  will answer by sending a  $\text{reply}()$  message containing the old value  $v_{old}$  and the previous value (no more valid) while  $s_j$  will answer by sending  $v_{old}$  and  $v'$ . However, when the client evaluates the  $\text{replies}$  set, it will find an occurrence of  $v_{old}$  for any alive server and evaluating the condition at line 12, Fig. 2(a) the reader will select and return  $v_{old}$  and we have a contradiction.

– **Case 2.3** -  $\mathbf{t_B(op')} + 2\delta < \mathbf{t_B(op)} < \mathbf{t_E(op')}$ .

In this case, the  $\text{READ}()$  message may be delivered also after the end of  $op'$ . If no more  $\text{write}()$  operations occur before the end of  $op$ , we fall down to case 2.1 and the claim follows. Contrarily, we fall down the the situation of case 2.2 where the concurrent write is a new write and the claim follow again.

□<sub>Lemma 10</sub>

**Theorem 2.** *Let  $c_\ell$  be a client invoking a  $\text{read}()$  operation  $op$ . If all the servers behave honestly, the protocol shown in Fig. 2 - 4 implements a regular register.*

**Proof** The proof directly follows by Theorem 1 and Lemma 10.

□<sub>Theorem 2</sub>

**Lemma 11.** *Let  $c_j$  be a client sending a  $\text{READ}()$  message at time  $t$  and let  $\text{REPLY}(< i, ts, v, ots, ov >)$  be the message delivered by  $c_j$  at time  $t' > t$  as reply to the  $\text{READ}()$ . Let  $lts$  be the value stored locally in the  $\text{last\_ts}$  variable by  $c_j$  at time  $t'$ . If  $|lts - ts| > 1$  then  $s_i$  is malicious.*

**Proof** Let us suppose by contradiction that  $s_i$  sends to  $c_j$  a  $\text{REPLY}(< i, ts, v, ots, ov >)$  such that  $|lts - ts| > 1$  and  $s_i$  is honest.

The  $\text{last\_ts}$  local variable is updated to a certain value  $x$  by any client  $c_j$  when (i)  $c_j$  delivered an  $\text{ACK}(< j, x, - >)$  message from any server it believes honest and (ii)  $x$  is strictly greater than the previous value stored in  $\text{last\_ts}$  (line 17, Fig. 3(a)).

Let  $t'$  be the time at which  $c_j$  delivers the  $\text{REPLY}(< i, ts, v, ots, ov >)$  message and let us denote with  $t_{up}$  the time at which  $c_j$  updated its  $\text{last\_ts}$  local variable to  $lts$ . If  $t_{ack}$  is the time at which  $s_i$  delivered such  $\text{WRITE}(v, lts)$  message and it sent back the  $\text{ACK}(< i, lts, - >)$  message to clients, it follows that  $t_{ack} < t_{up}$ .

If at time  $t'$   $lts \leq ts$ , it follows that  $t_{up} \leq t'$ , otherwise  $t_{up} > t'$ . Let us consider the two cases separately.

– **Case 1** -  $\mathbf{t_{up}} \leq \mathbf{t'}$ . In this case  $t' > t_{ack}$ . Let us note that the value  $ts$  sent by  $s_i$  in the  $\text{REPLY}$  message is the value stored locally by  $s_i$  in the variable  $ts_i$ . Let us call  $t_{rp}$  the time at which  $s_i$  sends the  $\text{REPLY}(< i, ts, v, ots, ov >)$  message to  $c_j$  and let us consider the following cases:

- **Case 1.1** -  $\mathbf{t_{rp}} < \mathbf{t_{ack}}$ . In this case,  $ts < lts$  and in particular, considering that  $\text{write}()$  operations are sequential, due to Corollary 1,  $ts = lts - 1$ . Thus,  $|lts - ts| = 1$  and we have a contradiction.

- **Case 1.2** -  $t_{rp} > t_{ack}$ . In this case,  $ts = lts$  and again we have a contradiction.
- **Case 2** -  $t_{up} > t'$ . Let us note that the value  $ts$  sent by  $s_i$  in the `REPLY` message is the value stored locally by  $s_i$  in the variable  $ts_i$ . Let us call  $t_{rp}$  the time at which  $s_i$  sends the `REPLY`( $\langle i, ts, v, ots, ov \rangle$ ) message to  $c_j$  and let us consider the following cases:
  - **Case 2.1** -  $t_{rp} < t_{ack}$ . In this case,  $ts = lts$  and we have a contradiction.
  - **Case 2.2** -  $t_{rp} > t_{ack}$ . In this case,  $ts < lts$  and in particular, considering that `write()` operations are sequential, due to Corollary 1,  $ts = lts - 1$ . Thus,  $|lts - ts| = 1$  and we have a contradiction.

□<sub>Lemma 11</sub>

**Lemma 12.** *Let  $c_j$  be a client sending a `READ()` message at time  $t$  and let `REPLY`( $\langle i, ts, v, ots, ov \rangle$ ) be the message delivered to  $c_j$  at time  $t' > t$  as reply to the `READ()`. Let  $mlts$  be the value stored locally in the `my_last_ts` variable and let  $mv$  be the value stored locally in the `my_last_value` variable by  $c_j$  at time  $t'$ . If  $mlts = ts$  and  $mv \neq v$  then  $s_i$  is malicious.*

**Proof** The proof simply follows by considering that a client  $c_j$  updates its `my_last_ts` and `my_last_val` local variables only at the beginning of a `write()` operation.

Due to Corollary 1, every `write()` operation has a unique timestamp thus if  $c_j$  delivers a value  $v \neq mv$  it means that the server altered it and it cannot be honest.

□<sub>Lemma 12</sub>

**Lemma 13.** *Let  $c_j$  be a client sending a `WRITE`( $\langle j, ts, val \rangle$ ) message at time  $t$ . If there exists a server  $s_j$  such that  $s_j$  does not send the triple  $\langle j, ts, val \rangle$  by time  $t + 3\delta$ , then  $s_j$  is faulty.*

**Proof** Considering that (i) the first `READ()` request is sent at time  $t + \delta$  (line 06, Fig. 3(a)), (ii) the `detection()` procedure on `replies` set is performed at time  $t + 2\delta$  (line 11, Fig. 3(a)) and (iii) there not exists concurrent `write()` operations, then any honest server will always provide the expected triple and the claim follows. □<sub>Lemma 13</sub>

**Lemma 14.** *Let  $c_j$  a client sending a `WRITE`( $\langle j, ts, val \rangle$ ) message at time  $t$ . If there exists a server  $s_j$  such that the triple  $\langle j, ts, - \rangle$  does not appear in the `ack` set at time  $t + 2\delta$ , then  $s_j$  is detected as faulty.*

**Proof** Considering that (i) the `detection()` procedure on `ack` set is performed at time  $t + 2\delta$  (line 09, Fig. 3(a)) and (ii) there not exists concurrent `write()` operations, then a honest server will provide the expected triple and the claim follows. □<sub>Lemma 14</sub>

**Theorem 3.** *A honest alive server  $s_i$  is never detected as malicious.*

**Proof** A server  $s_i$  deviates from the protocol if:

1.  $s_i$  omits to send  $\text{ACK}()$  messages during a  $\text{write}()$  or to sends  $\text{REPLY}()$  messages during a  $\text{read}()$  (lines 02-07, Fig. 4).
2.  $s_i$  sends bad timestamps (i.e. too old, lines 19-23, Fig. 4 or too new, lines 32-36, Fig. 4).
3.  $s_i$  sends a pair  $\langle \text{value}, ts \rangle$  with the correct timestamp and the wrong value (lines 24-31, Fig. 4).

Thus, a honest server may be erroneously detected as faulty only if one of the previous cases occur. However, due to Lemmas 11 - 14, we have that if one of the previous situation happens,  $s_i$  is necessarily malicious and the claim follows.

$\square_{\text{Theorem 3}}$

**Theorem 4.** *If  $D_s, G_s$  and  $\theta$  are such that equation (4) holds then the protocol shown in Fig. 2 - 4 implements a regular register with only 1 alive honest server.*

**Proof** Let us first notice that the existence of the detection function creates the dualism between risky requests and risk-less request. In addition, the existence of one honest alive server prevents the attacker to create a collusion able to make the condition at lines 08, 12 and 17 satisfied with not valid values. Thus, the claim simply follows by Lemma 3 and Theorem 2.

$\square_{\text{Theorem 4}}$

## 7 $\mathcal{P}_{cv}$ and $\mathcal{P}_{hash}$ Protocols for a Regular Register when $D_s \geq G_s$

In the following, we show how to modify the protocol to get  $\theta \geq \frac{1}{2}$ , when  $D_s \geq G_s$ . In particular, we propose two possible extensions: the first using a probabilistic collaborative detection at the client side (introducing a cost in terms of number of messages needed to run the detection) and the second using a kind of fingerprint to prevent servers misbehavior (introducing a computational cost).

**A collaborative detection protocol  $\mathcal{P}_{cv}$ .** The collaborative detection involves all the clients in the detection process and exploits the fact that the last writer remains in the system and it is always able to identify a faulty server. The basic idea is to look for a write witness (i.e., the writer) each time that a reader is not able to decide about the correctness of a value. This solution allows to identify malicious server and to decide and return always a correct value. However, considering that (i) we want to decouple as much as possible servers and client, (ii) this collaborative approach has a cost in terms of messages and (iii) to force rational servers to behave correctly it is sufficient to get  $\theta \geq \frac{1}{2}$  (according to Lemma 3), then we use this collaborative approach only with a given probability. More in details, in  $\mathcal{P}_{cv}$  protocol, when a reader does not collect the same value from all servers it flips a coin to decide if running the collaborative detection or not. If the outcome is 1, then it broadcasts to all the other clients the timestamps collected during the read operation and waits that some writer acknowledge them. When a client receives a check timestamp request, it checks if it corresponds to its last written value and if so, it replies with such a value so that the reader can double-check information provided by servers. If there is no match between values and timestamps, then clients are able to detect a faulty server and exclude it from the computation.

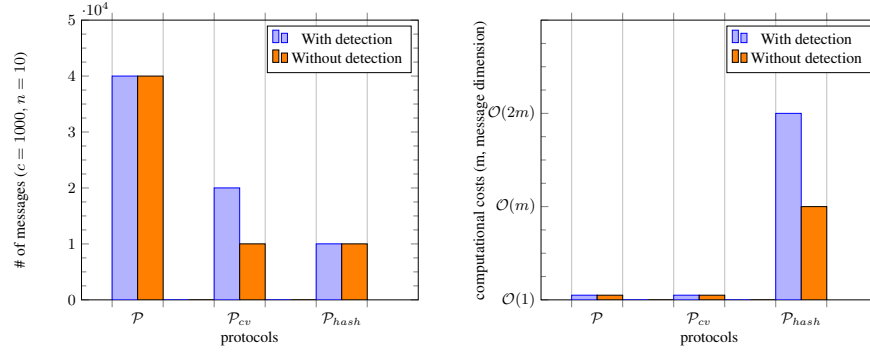
The introduction of this probabilistic step in the protocol increases the value of  $\theta$  to  $\frac{1}{2}$ . As a consequence, following Lemma 3, any rational server will decide to behave correctly to avoid to be detected.

**A fingerprint-based detection protocol  $\mathcal{P}_{hash}$ .** Let us recall that the basic idea behind the detection process is to include inside reply messages (i.e., write acknowledgements or read replies) “enough” information to verify the correctness of the provided information. In particular, in protocol  $\mathcal{P}$ , servers are required to acknowledge write operations by sending back the corresponding timestamp so that each client is always aware about it and the writer is able to verify that no bad timestamps are sent to clients.

In protocol  $\mathcal{P}_{hash}$ , the basic idea is to extend  $\mathcal{P}$  by including another information i.e., a fingerprint of the value and its timestamp (e.g., its hash), in the write message and in its acknowledgement so that it is always possible for a client to check that servers are replying correctly. More in details, when a client writes, it computes the hash of the value and its corresponding timestamp and attaches such fingerprint to the message. In such way (as for  $\mathcal{P}$ ) when servers acknowledge a write, they send back the correct fingerprint to all clients. Having such information, all clients are potentially able to detect locally if values collected during a read operation are never written values (this can be simply done by computing the hash of the message and compare it with the one received during the last write). However, as in the case of  $\mathcal{P}_{cv}$ , this detection has a cost and, to get  $\theta \geq \frac{1}{2}$  it is sufficient that this is done with a certain probability. Thus, when a reader does not collect the same value from all servers, it flips a coin and if the outcome is 1 then it computes the hash of the messages it delivered and compares them with the hashes it knows to be associated to a specific timestamp. The introduction of this step is enough to get  $\theta = \frac{1}{2}$  and to prevent rational servers deviating from the protocol. Notice that, as for  $\mathcal{P}_{cv}$ , the employment of the random coin has a twofold purpose: (i) to provide a solution for  $D_s \geq G_s$ , for which it is enough to have  $\theta \geq \frac{1}{2}$  and (ii) to avoid to always perform the costly detection operation.

**Trade offs.** Figure 5 shows a qualitative comparison of the three proposed protocols in terms of message complexity and computational cost. In particular, we compare the cost of the protocols both in presence and absence of a server attack (i.e., when the detection is necessary or not). As we can see,  $\mathcal{P}$  requires the highest number of messages and such number does not depend on the real need of doing detection but it is rather required to mask the type of operation that a client is doing and to make indistinguishable real read messages from dummy ones. Concerning its computational cost, it is constant since it does not depend on the message size.

In  $\mathcal{P}_{cv}$  it is possible to save the dummy read messages as we do not need anymore to mask the message pattern but we need to pay the cost of the collaborative detection, if it is needed. In fact, if a reader is not able to decide a value, it needs to send messages to contact all the other clients (higher message complexity in case of server misbehaviour). Concerning the computational cost, it is not affected by the detection. Conversely,  $\mathcal{P}_{hash}$  exhibits the dual behaviour: message complexity is not affected by server misbehaviour but the computational cost is impacted by the need of detection.



**Fig. 5.** Qualitative analysis of protocols with respect to their message complexity (left figure) and computational complexity (right figure). For the message complexity we consider a system where the number of servers is  $n = 10$  and the number of clients is  $c = 1000$ . For the computational complexity we consider the cost with respect to the message size  $m$ .

Thus, we can conclude saying that  $\mathcal{P}$  is a pessimistic protocol and is the most expensive one but it allows to maintain clients and servers completely decoupled. Contrarily,  $\mathcal{P}_{cv}$  and  $\mathcal{P}_{hash}$  are optimistic as they perform lightweight operations and, if needed, they perform an heavy detection (with a high message cost in the case of  $\mathcal{P}_{cv}$  and a high computational cost in case of  $\mathcal{P}_{hash}$ ).

#### Sketch of $\mathcal{P}_{cv}$ and $\mathcal{P}_{hash}$ correctness

$\mathcal{P}_{cv}$  has not variations with respect to  $\mathcal{P}$  concerning the write and read operations (the detection procedure is removed from the write operation). Thus the algorithm is proved to work. Concerning the detection part we have that when a client fails to read, then it asks to other clients to reach the last writing one that replies with the correct value. Thus the reading client is able to detect servers acting in a malicious way. Since the detection procedure is run depending on the outcome of a random bit then  $\theta = \frac{1}{2}$ . Thus if  $D_s > G_s$  then (4) holds then  $\mathcal{P}_{cv}$  implements a regular register with only 1 alive honest server.

$\mathcal{P}_{hash}$  has not variations with respect to  $\mathcal{P}$  concerning the write and read operations (the detection procedure is removed from the write operation). Thus the algorithm is proved to work. The difference with  $\mathcal{P}$  is that the detection procedure slightly modifies the write operation. Each time a new value is written, the hash is computed and spread along with the timestamp. The proof of that directly follows generalizing Lemma 6. When a client fails to read then it performs the hash of the delivered replies to compare them to the one it knows to be associated to the related timestamp. Thus the reading client is able to detect servers acting in a malicious way. Since the detection procedure is run depending on the outcome of a random bit then  $\theta = \frac{1}{2}$ , if  $D_s > G_s$  then (4) holds then  $\mathcal{P}_{cv}$  implements a regular register with only 1 alive honest server.

## 8 Conclusion

This paper addresses the problem of building a regular register in a distributed system where clients are anonymous and servers maintaining the register state may be rational malicious processes. We have modelled our problem as a two-parties Bayesian game and we designed distributed protocols able to reach the Bayesian Nash Equilibrium and to emulate a regular register when the loss in case of detection is greater than the gain obtained from the deviation (i.e.  $D_s > G_s$ ). To the best of our knowledge, our protocols are the first register protocols working in the absence of knowledge on the number of compromised replicas.

The protocols rely on the following assumptions: (i) rational malicious servers act independently and do not form a coalition, (ii) the system is synchronous, (iii) clients are anonymous and (iv) write operations are serialised.

As future works, we are investigating how to solve the same problem under weaker synchrony assumption or in the case an attacker controls a coalition of processes. Addressing these points is actually far from be trivial. Considering a fully asynchronous system, in fact, makes impossible to use our punishment mechanism as clients are not able to distinguish alive but silent servers from those crashed. Additionally, when the attacker is able to compromise and control a coalition of processes, the model provided in this paper is no more adequate and we are studying if and how it is possible to define a *Bayesian Coalitional Game* [12] for our problem and if an equilibrium can be reached in this case.

## Acknowledgments

This present work has been partially supported by the EURASIA project and CINI Cybersecurity National Laboratory within the project FilieraSicura: Securing the Supply Chain of Domestic Critical Infrastructures from Cyber Attacks ([www.filierasicura.it](http://www.filierasicura.it)) funded by CISCO Systems Inc. and Leonardo SpA.

## References

1. Abraham, I., Dolev, D., Halpern, J. Y. *Distributed Protocols for Leader Election: A Game-Theoretic Perspective*. DISC 2013: 61-75
2. Afek, Y., Ginzberg, Y., Landau Feibish, S. and Sulamy, M. *Distributed computing building blocks for rational agents*. PODC 2014: 406-415.
3. Aiyer, A. S., Alvisi, L., Clement, A., Dahlin, M., Martin, J. P., and Porth, C. *BAR fault tolerance for cooperative services*. ACM SIGOPS Operating Systems Review. ACM, 2005. p. 45-58.
4. Attiya, H., Bar-Noy, A., and Dolev, D. *Sharing memory robustly in message-passing systems*. Journal of the ACM 42, 1, 1995, 124-142.
5. Bazzi R. A., *Synchronous Byzantine Quorum Systems*, Distributed Computing 13(1), 45-52, 2000.
6. Chaudhuri S., Kosa M.J. and Welch J., *One-write Algorithms for Multivalued Regular and Atomic Registers*. Acta Informatica, 37:161-192, 2000.
7. Clement, A., Li, H. C., Napper, J., Martin, J., Alvisi, L., Dahlin, M. *BAR primer*, DSN 2008: 287-296
8. Clement, A., Napper, J., Li, H., Martin, J. P., Alvisi, L., and Dahlin, M. *Theory of BAR games*, PODC 2007: 358-359
9. Delporte-Gallet, C., Fauconnier, H., Tran-The, H. *Uniform Consensus with Homonyms and Omission Failures* ICDCN 2013: 161-175
10. Fudenberg, D., Tirole, J. *Game theory*, 1991. Cambridge, Massachusetts.
11. Haldar S. and Vidyasankar K., *Constructing 1-writer Multireader Multivalued Atomic Variables from Regular Variables*. JACM, 42(1):186-203, 1995.
12. Jeong, S., Shoham, Y. *Bayesian Coalitional Games*. AAAI. 2008: 95-100
13. Li, H. C., Clement, A., Wong, E. L., Napper, J., Roy, I., Alvisi, L., Dahlin, M. *BAR Gossip* OSDI 2006: 191-204
14. Lamport. L., On Interprocess Communication, Part 1: Models, Part 2: Algorithms, *Distributed Computing*, 1(2):77-101, 1986.
15. Mahajan, P., Setty, S., Lee, S., Clement, A., Alvisi, L., Dahlin, M., and Walfish, M. *Depot: Cloud storage with minimal trust*, ACM TOCS 29(4), 2011
16. Malkhi D., Reiter M. K. *Byzantine Quorum Systems*, Distributed Computing 11(4), 203-213, 1998.
17. Martin J., Alvisi L., Dahlin M., *Small Byzantine Quorum Systems*, DSN 2002: 374-388.
18. Martin J., Alvisi L., Dahlin M.. *Minimal Byzantine Storage*, DISC 2002.
19. Schneider, F. B. *Implementing fault-tolerant services using the state machine approach: A tutorial*, ACM Computing Surveys 22(4): 299-319, 1990.
20. Singh A.K., Anderson J.H. and Gouda M., *The Elusive Atomic Register*. JACM, 41(2):331-334, 1994.
21. Sousa, P., Bessani, A. N., Correia, M., Neves, N. F., Verissimo, P. *Highly available intrusion-tolerant services with proactive-reactive recovery*, IEEE TPDS 21(4): 452-465, 2010 .
22. The Tor Project <https://www.torproject.org>.
23. Vidyasankar K., *Converting Lamport's Regular Register to Atomic Register*. IPL, 28(6):287-290, 1988
24. Vityani P. and Awerbuch B., *Atomic Shared Register Access by Asynchronous Hardware*. FOCS 1987, 223-243.
25. Ostrovsky, R. and Yung, M., *How to withstand mobile virus attacks*. PODC 1991, 51-59.
26. Shamir, A. *How to share a secret* Comm. of ACM 1979, 612-613.
27. Dolev, S., ElDefrawy, K., Lampkins, J., Ostrovsky, R., and Yung, M. *Proactive secret sharing with a dishonest majority*. SCN 2016, 529-548.
28. Cramer, R. and Damgard, I. B. *Secure Multiparty Computation* Cambridge University Press 2015.